WHAT IS CLAIMED IS:

1. A cluster comprising:

a plurality of filesystems, each of the plurality of filesystems included in one of a plurality of service groups, and each of the plurality of service groups including one or more filesystems of the plurality of filesystems; and

a plurality of nodes configured to: (i) act as a server for the plurality of filesystems; (ii) provide record locking services in the plurality of filesystems; and (iii) maintain a plurality of client lists, each of the plurality of client lists included in a respective service group of the plurality of service groups and identifying clients having at least one lock on one of the one or more filesystems included in the respective service group.

2. The cluster as recited in claim 1 wherein each of the plurality of client lists is stored in at least one of the one or more filesystems included in the respective service group.

3. The cluster as recited in claim 1 wherein the plurality of nodes are configured to fail over a first service group of the plurality of service groups from a first node of the plurality of nodes to a second node of the plurality of nodes, and wherein the second node is configured to initiate lock recovery for locks in the one or more filesystems included in the first service group responsive to the fail over.

4. The cluster as recited in claim 3 wherein the second node is configured to read a first client list of the plurality of client lists, the first client list included in the first service group, to initiate lock recovery.

5. The cluster as recited in claim 3 wherein the second node is configured to initiate lock recovery by notifying each of the clients in the first client list that the clients should reclaim locks previously granted on a filesystem within the first service group.

6. The cluster as recited in claim 5 wherein the second node is configured to notify each of the clients by transmitting one or more server identifiers within the first service group to the clients.

7. The cluster as recited in claim 6 wherein each of the one or more server identifiers comprises an internet protocol address.

8. The cluster as recited in claim 5 wherein the second node is configured to initiate a period of time for the clients in the first client list to reclaim locks, wherein the second node is configured not to grant new locks on filesystems within the first service group during the period.

9. The cluster as recited in claim 8 wherein, if the second node is also acting as a server for one or more filesystems in a second service group of the plurality of service groups, the second node is configured to grant a lock requested by a client in a filesystem within the second service group during the period.

10. The cluster as recited in claim 3 wherein, if the second node is also acting as a server for one or more filesystems in a second service group of the plurality of service groups, the second node is configured to not initiate lock recovery for locks in the second service group responsive to the fail over of the first service group.

11. A method comprising:

maintaining a plurality of client lists, each of the plurality of client lists included

in a respective service group of a plurality of service groups and identifying clients having at least one lock on at least one filesystem included in the respective service group;

failing over a first service group of the plurality of service groups from a first node of a plurality of nodes to a second node of the plurality of nodes; and

the second node initiating lock recovery for locks on one or more filesystems including in the first service group responsive to the fail over using a first client list of the plurality of client lists, the first client list included in the first service group.

12. The method as recited in claim 11 further comprising storing each of the plurality of client lists in at least one of the one or more filesystems included in the respective service group.

13. The method as recited in claim 11 wherein initiating lock recovery comprises notifying each of the clients in the first client list that the clients should reclaim locks previously granted on at least one filesystem in the first service group.

14. The method as recited in claim 13 wherein notifying each of the clients comprises transmitting one or more server identifiers included in the first service group.

15. The method as recited in claim 13 further comprising:

initiating a period of time for the clients in the first client list to reclaim locks in the filesystems included in the first service group; and

the second node not granting new locks in the filesystems included in the first

service group during the period.

16. The method as recited in claim 15 wherein the second node is also acting as a server for one or more filesystems in a second service group of the plurality of service groups, the method further comprising the second node not interrupting locking services for the second service group during the period.

17. The method as recited in claim 16 wherein not interrupting locking services comprises granting a lock in a filesystem within the second service group during the period.

18. The method as recited in claim 15 wherein the first node is also acting as a server for one or more filesystems in a second service group of the plurality of service groups, the method further comprising the first node not interrupting locking services for the second service group during the period.

19. The method as recited in claim 11 wherein the second node is also acting as a server for one or more filesystems in a second service group of the plurality of service groups, the method further comprising the second node not interrupting locking services for the second service group during the period.

20. A computer accessible medium encoded with a plurality of instructions which, when executed in a first node of a plurality of nodes in response to a fail over of a first service group of a plurality of service groups from a second node of the plurality of nodes and each of the plurality of service groups including at least one filesystem, initiate lock recovery for locks on each filesystem in the first service group using a first client list of a plurality of client lists, wherein each of the plurality of client lists is included in a respective service group of the plurality of service groups and identifies clients having at least one lock in at least one filesystem in the respective service group, and wherein the

first client list is included in the first service group.

21. The computer accessible medium as recited in claim 20 wherein each of the plurality of client lists is stored in the respective service group.

22. The computer accessible medium as recited in claim 20 wherein the plurality of instructions, when executed, initiate lock recovery by notifying each of the clients in the first client list that the clients should reclaim locks previously granted on a filesystem in the first service group.

23. The computer accessible medium as recited in claim 22 wherein notifying each of the clients in the first client lists comprises transmitting one or more server identifiers included in the first service group.

24. The computer accessible medium as recited in claim 22 wherein the plurality of instructions, when executed:

initiate a period of time for the clients in the first client list to reclaim locks on the at least one filesystem in the first service group; and

do not grant new locks on the at least one filesystem in the first service group during the period.

25. The computer accessible medium as recited in claim 24 wherein, if the second node is also acting as a server for one or more filesystems in a second service group of the plurality of service groups, the plurality of instructions, when executed, grant a lock requested by a client on a filesystem in the second service group during the period.

26. The computer accessible medium as recited in claim 20 wherein, if the second node

is also acting as a server for one or more filesystems in a second service group of the plurality of service groups, the plurality of instructions, when executed, do not initiate lock recovery for locks on filesystems in the second service group responsive to the fail over of the first service group.

27. A computer accessible medium encoded with a plurality of instructions which, when executed in a first node of a plurality of nodes in response to a fail over of a first service group of a plurality of service groups from a second node of the plurality of nodes and each of the plurality of service groups comprises at least one filesystem, initiate lock recovery for locks on the at least one filesystem in the first service group, and wherein locks are maintained in the at least one filesystem in a second service group of the plurality of service groups during a time period that locks in the first service group are recovered.

28. The computer accessible medium as recited in claim 27 wherein the first node is acting as a server for the at least one filesystem in the second service group.

29. The computer accessible medium as recited in claim 27 wherein a different node of the plurality of nodes is acting as the server for the at least one filesystem in the second service group.

30. The computer accessible medium as recited in claim 27 wherein the plurality of instructions, when executed:

initiate a period of time for clients to reclaim locks on the at least one filesystem in the first service group; and

do not grant new locks on the at least one filesystem in the first service group during the period.

31. The computer accessible medium as recited in claim 30 wherein the plurality of instructions, when executed, grant a lock requested by a client on one of the at least one filesystems included in the second service group during the period.

32. The computer accessible medium as recited in claim 27 wherein the plurality of instructions, when executed, do not interrupt locking services for the second service group.

33. A method comprising:

initiating lock recovery for locks on one or more filesystems in a first service group of a plurality of service groups in response to failing over the first service group to a first node of a plurality of nodes from a second node of the plurality of nodes; and

maintaining locks on one or more filesystems in a second service group of the plurality of service groups during a time period that locks in the first service group are recovered.

34. The method as recited in claim 33 wherein the first node is acting as a server for the one or more filesystems in the second service group.

35. The method as recited in claim 33 wherein a different node of the plurality of nodes is acting as the server for the one or more filesystems in the second service group.

36. The method as recited in claim 33 further comprising:

initiating a period of time for clients to reclaim locks on the one or more

filesystems in the first service group; and

not granting new locks on the one or more filesystems in the first service group during the period.

37. The method as recited in claim 36 further comprising granting a lock requested by a client on a filesystem in the second service group during the period.